

ABSTRACT

A method and mechanism for pipelined table functions is disclosed. With pipelining, data from a first table function does not have to be materialized into a complete collection before it is consumed by a succeeding table function. Instead, a producer of data creates a stream of data that is immediately utilized by a consumer of that data. Also disclosed is a method and mechanism for parallel processing of table functions, in which the set of work operated upon by a table function is sub-divided into smaller portions that are assigned to a plurality of table function slaves. Also disclosed is an integration between pipelining and parallelized execution for table functions.